

CLAIMS:

1. A method for maintaining aggregations of values contained by fields of multiple database records, comprising:

creating multiple aggregation groups, each group including a plurality of aggregation records, each aggregation record including a value for an aggregation of values contained by fields of a distinct subset of the multiple database records;

selecting a first aggregation group upon insertion or update of a first of the multiple database records;

revising, based on one or more values within the inserted or updated first database record and as part of a first aggregation group update transaction, the aggregation value of one of the aggregation records of the first aggregation group;

preventing subsequent selection of the first aggregation group until the first aggregation group update transaction is completed;

selecting, while the first aggregation group update transaction is being performed and upon insertion or update of a second of the multiple database records, a second aggregation group; and

revising, based on one or more values within the inserted or updated second database record and during the first aggregation group update transaction, the aggregation value of one of the aggregation records of the second aggregation group.

2. The method of claim 1, wherein each aggregation group is a separate partition of a multi-partition aggregation table.

3. The method of claim 1, further comprising:
selecting a third aggregation group upon initiation of a subsequent transaction to update the first of the multiple database records; and
revising, based on one or more values within the subsequently updated first database record, the aggregation value of an aggregation record of the third aggregation group.
4. The method of claim 1, further comprising:
combining the multiple aggregation groups into a single table of aggregation records, each record of the single table aggregating values of an aggregation record from each of the multiple aggregation groups.
5. The method of claim 1, wherein said creating multiple aggregation groups comprises creating at least three aggregation groups.
6. The method of claim 1, wherein said creating multiple aggregation groups comprises creating at least ten aggregation groups.
7. The method of claim 1, wherein said creating multiple aggregation groups comprises creating a number of aggregation groups, and wherein said number exceeds a number of processors in a computer used to perform the steps of claim 1.
8. The method of claim 1, wherein:

said selecting a first aggregation group is initiated in response to a request by a first program thread to access one of the multiple aggregation groups; and

said preventing subsequent selection comprises refusing other program threads access to the first aggregation group until the first program thread no longer requires access to the first aggregation group.

9. The method of claim 8, further comprising:

determining, upon receiving a request from a program thread for access to an aggregation group, a system identifier for the requesting program thread; and

assigning an aggregation group identifier to the requesting program thread based on the determined system identifier.

10. The method of claim 1, wherein:

each of the multiple database records corresponds to an instance of an organizational activity,

each of the multiple database records includes a field having a value indicating the corresponding instance to be in one of several process states, and

each aggregation group includes time-sorted aggregation records, each time-sorted aggregation record containing an aggregation value for instances in one of the several process states during a time period associated with the time-sorted aggregation record.

11. The method of claim 10, wherein one of said several process states comprises an instance being completed, and further comprising:

deleting aggregation records corresponding to instances completed outside of a preselected time window.

12. A method for maintaining aggregated data regarding multiple instances of an organizational activity, each instance of the activity having one of a plurality of process states, comprising:

creating a plurality of records in an aggregated data table, each record containing an aggregated value for a subset of the multiple instances in the same process state during the same time period; and

updating the aggregated data table to reflect deletion of data corresponding to instances in one of the process states outside of a preselected time window.

13. The method of claim 12, comprising:

deleting records from the aggregated data table indicating none of the multiple instances were in one of the process states during a time period corresponding to the deleted record.

14. The method of claim 12, wherein said updating the aggregated data table comprises deleting records containing aggregated data for instances of the multiple instances completed outside of the preselected time window.

15. The method of claim 12, wherein the one of the process states corresponds to an instance being completed and wherein each of the plurality of records includes a completion time field; and further comprising:

assigning a null value to the completion time field for records corresponding to instances that are not completed;

assigning a non-null value to the completion time field for records corresponding to instances in the completed process state.

16. The method of claim 15, wherein data for each instance is maintained in an associated record of an instances data table, and comprising:

determining if a record of the instances data table has been updated;

revising, upon determining that the instances data table record has been updated and based on said update, an aggregated value of one of the plurality of records;

further determining whether the updated instances data table record contains a value indicating the associated instance is in the completed process state;

generating, upon determining that the associated instance is in the completed process state, a time of completion value for the instance; and

updating a record of the aggregated data table based on the time of completion value for the instance.

17. The method of claim 16, comprising:

deleting records containing aggregated data for instances of the multiple instances completed outside of the preselected time window.

18. A computer-readable medium having stored thereon data representing sequences of instructions which, when executed by a processor, cause the processor to perform steps comprising:

creating multiple aggregation groups, each group including a plurality of aggregation records, each aggregation record including a value for an aggregation of values contained by fields of a distinct subset of multiple database records;

selecting a first aggregation group upon insertion or update of a first of the multiple database records;

revising, based on one or more values within the inserted or updated first database record and as part of a first aggregation group update transaction, the aggregation value of one of the aggregation records of the first aggregation group;

preventing subsequent selection of the first aggregation group until the first aggregation group update transaction is completed;

selecting, while the first aggregation group update transaction is being performed and upon insertion or update of a second of the multiple database records, a second aggregation group; and

revising, based on one or more values within the inserted or updated second database record and during the first aggregation group update transaction, the aggregation value of one of the aggregation records of the second aggregation group.

19. The computer-readable medium of claim 18, wherein each aggregation group is a separate partition of a multi-partition aggregation table.

20. The computer-readable medium of claim 18, comprising further instructions for performing steps comprising:

selecting a third aggregation group upon initiation of a subsequent transaction to update the first of the multiple database records;

revising, based on one or more values within the subsequently updated first database record, the aggregation value of an aggregation record of the third aggregation group.

21. The computer-readable medium of claim 18, comprising further instructions for performing steps comprising:

combining the multiple aggregation groups into a single table of aggregation records, each record of the single table aggregating values of an aggregation record from each of the multiple aggregation groups.

22. The computer-readable medium of claim 18, wherein said creating multiple aggregation groups comprises creating at least three aggregation groups.

23. The computer-readable medium of claim 18, wherein said creating multiple aggregation groups comprises creating at least ten aggregation groups.

24. The computer-readable medium of claim 18, wherein:

said selecting a first aggregation group is initiated in response to a request by a first program thread to access one of the multiple aggregation groups; and

said preventing subsequent selection comprises refusing other program threads access to the first aggregation group until the first program thread no longer requires access to the first aggregation group.

25. The computer-readable medium of claim 24, comprising further instructions for performing steps comprising:

determining, upon receiving a request from a program thread for access to an aggregation group, a system identifier for the requesting program thread; and

assigning an aggregation group identifier to the requesting program thread based on the determined system identifier.

26. The computer-readable medium of claim 18, wherein:

each of the multiple database records corresponds to an instance of an organizational activity,

each of the multiple database records includes a field having a value indicating the corresponding instance to be in one of several process states, and

each aggregation group includes time-sorted aggregation records, each time-sorted aggregation record containing an aggregation value for instances in one of the several process states during a time period associated with the time-sorted aggregation record.

27. The computer-readable medium of claim 26, wherein one of said several process states comprises an instance being completed, and comprising further instructions for performing steps comprising:

deleting aggregation records corresponding to instances completed outside of a preselected time window.

28. A computer-readable medium having stored thereon data representing sequences of instructions which, when executed by a processor, cause the processor to perform steps comprising:

creating a plurality of records in an aggregated data table, each record containing an aggregated value for a subset of multiple instances of an organizational activity, each instance of the activity having one of a plurality of process states, each instance in a subset being in the same process state during the same time period; and

updating the aggregated data table to reflect deletion of data corresponding to instances in one of the process states outside of a preselected time window.

29. The computer-readable medium of claim 28, comprising further instructions for performing steps comprising:

deleting records from the aggregated data table indicating none of the multiple instances were in one of the process states during a time period corresponding to the deleted record.

30. The computer-readable medium of claim 28, wherein said updating the aggregated data table comprises deleting records containing aggregated data for instances of the multiple instances completed outside of the preselected time window.

31. The computer-readable medium of claim 28, wherein the one of the process states corresponds to an instance being completed and wherein each of the plurality of records includes a completion time field; and comprising further instructions for performing steps comprising:

assigning a null value to the completion time field for records corresponding to instances that are not completed;

assigning a non-null value to the completion time field for records corresponding to instances in the completed process state.

32. The computer-readable medium of claim 28, wherein data for each instance is maintained in an associated record of an instances data table, and comprising further instructions for performing steps comprising:

determining if a record of the instances data table has been updated;

revising, upon determining that the instances data table record has been updated and based on said update, an aggregated value of one of the plurality of records;

further determining whether the updated instances data table record contains a value indicating the associated instance is in the completed process state;

generating, upon determining that the associated instance is in the completed process state, a time of completion value for the instance; and

updating a record of the aggregated data table based on the time of completion value for the instance.

33. The computer-readable medium of claim 32, comprising further instructions for performing steps comprising:

deleting records containing aggregated data for instances of the multiple instances completed outside of the preselected time window.

34. A data processing apparatus for maintaining aggregations of values contained by fields of multiple database records, comprising:

at least one data storage device;

at least one user input device; and

a processor operatively connected to said storage device and said user input device, wherein the at least one data storage device has stored thereon a set of instructions which, when executed, configure said processor to:

create multiple aggregation groups, each group including a plurality of aggregation records, each aggregation record including a value for an aggregation of values contained by fields of a distinct subset of multiple database records,

select a first aggregation group upon insertion or update of a first of the multiple database records,

revise, based on one or more values within the inserted or updated first database record and as part of a first aggregation group update transaction, the aggregation value of one of the aggregation records of the first aggregation group,

prevent subsequent selection of the first aggregation group until the first aggregation group update transaction is completed,

select, while the first aggregation group update transaction is being performed and upon insertion or update of a second of the multiple database records, a second aggregation group, and

revise, based on one or more values within the inserted or updated second database record and during the first aggregation group update transaction, the aggregation value of one of the aggregation records of the second aggregation group.

35. The data processing apparatus of claim 34, wherein the set of instructions includes additional instructions which, when executed, configure said processor to:

select a third aggregation group upon initiation of a subsequent transaction to update the first of the multiple database records, and

revise, based on one or more values within the subsequently updated first database record, the aggregation value of an aggregation record of the third aggregation group.

36. The data processing apparatus of claim 34, wherein the set of instructions includes additional instructions which, when executed, configure said processor to:

combine the multiple aggregation groups into a single table of aggregation records, each record of the single table aggregating values of an aggregation record from each of the multiple aggregation groups.

37. The data processing apparatus of claim 34, wherein

said selecting a first aggregation group is initiated in response to a request by a first program thread to access one of the multiple aggregation groups, and

said preventing subsequent selection comprises refusing other program threads access to the first aggregation group until the first program thread no longer requires access to the first aggregation group.

38. The data processing apparatus of claim 37, wherein the set of instructions includes additional instructions which, when executed, configure said processor to:

determine, upon receiving a request from a program thread for access to an aggregation group, a system identifier for the requesting program thread, and

assign an aggregation group identifier to the requesting program thread based on the determined system identifier.

39. A data processing apparatus for maintaining aggregated data regarding multiple instances of an organizational activity, each instance of the activity having one of a plurality of process states, comprising:

at least one data storage device;

at least one user input device; and

a processor operatively connected to said storage device and said user input device, wherein the at least one data storage device has stored thereon a set of instructions which, when executed, configure said processor to:

create a plurality of records in an aggregated data table, each record containing an aggregated value for a subset of the multiple instances in the same process state during the same time period, and

update the aggregated data table to reflect deletion of data corresponding to instances in one of the process states outside of a preselected time window.

40. The data processing apparatus of claim 39, wherein the set of instructions includes additional instructions which, when executed, configure said processor to:

delete records from the aggregated data table indicating none of the multiple instances were in one of the process states during a time period corresponding to the deleted record.

41. The data processing apparatus of claim 39, wherein said updating the aggregated data table comprises deleting records containing aggregated data for instances of the multiple instances completed outside of the preselected time window.

42. The data processing apparatus of claim 39, wherein the one of the process states corresponds to an instance being completed, wherein each of the plurality of records includes a completion time field, and wherein the set of instructions includes additional instructions which, when executed, configure said processor to:

assign a null value to the completion time field for records corresponding to instances that are not completed, and

assign a non-null value to the completion time field for records corresponding to instances in the completed process state.

43. The data processing apparatus of claim 42, wherein data for each instance is maintained in an associated record of an instances data table, and wherein the set of instructions includes additional instructions which, when executed, configure said processor to:

determine if a record of the instances data table has been updated,
revise, upon determining that the instances data table record has been updated and based on said update, an aggregated value of one of the plurality of records,
further determine whether the updated instances data table record contains a value indicating the associated instance is in the completed process state,
generate, upon determining that the associated instance is in the completed process state, a time of completion value for the instance, and
update a record of the aggregated data table based on the time of completion value for the instance.

44. The data processing apparatus of claim 43, wherein the set of instructions includes additional instructions which, when executed, configure said processor to:

delete records containing aggregated data for instances of the multiple instances completed outside of the preselected time window.